
Uniform Deviation Bounds for k -Means Clustering

Olivier Bachem¹ Mario Lucic¹ S. Hamed Hassani¹ Andreas Krause¹

Abstract

Uniform deviation bounds limit the difference between a model’s expected loss and its loss on a random sample *uniformly* for all models in a learning problem. In this paper, we provide a novel framework to obtain uniform deviation bounds for *unbounded* loss functions. As a result, we obtain competitive uniform deviation bounds for k -Means clustering under weak assumptions on the underlying distribution. If the fourth moment is bounded, we prove a rate of $\mathcal{O}\left(m^{-\frac{1}{2}}\right)$ compared to the previously known $\mathcal{O}\left(m^{-\frac{1}{4}}\right)$ rate. We further show that this rate also depends on the kurtosis — the normalized fourth moment which measures the “tailedness” of the distribution. We also provide improved rates under progressively stronger assumptions, namely, bounded higher moments, subgaussianity and bounded support of the underlying distribution.

1. Introduction

Empirical risk minimization — i.e. the training of models on a finite sample drawn i.i.d from an underlying distribution — is a central paradigm in machine learning. The hope is that models trained on the finite sample perform provably well even on previously unseen samples from the underlying distribution. But how many samples m are required to guarantee a low approximation error ϵ ? *Uniform deviation bounds* provide the answer. Informally, they are the worst-case difference across all possible models between the empirical loss of a model and its expected loss. As such, they determine how many samples are required to achieve a fixed error in terms of the loss function. In this paper, we consider the popular k -Means clustering problem and provide uniform deviation bounds based on *weak* assumptions on the underlying data generating distribution.

¹Department of Computer Science, ETH Zurich. Correspondence to: Olivier Bachem <olivier.bachem@inf.ethz.ch>.

Related work. Traditional *Vapnik-Chervonenkis* theory provides tools to obtain uniform deviation bounds for binary concept classes such as classification using halfspaces (Vapnik & Chervonenkis, 1971). While these results have been extended to provide uniform deviation bounds for sets of continuous functions bounded in $[0, 1]$ (Haussler, 1992; Li et al., 2001), these results are not readily applied to k -Means clustering as the underlying loss function in k -Means clustering is continuous and unbounded.

In their seminal work, Pollard (1981) shows that k -Means clustering is *strongly consistent*, i.e., that the optimal cluster centers on a random sample converge almost surely to the optimal centers of the distribution under a weak assumption. This has sparked a long line of research on *cluster stability* (Ben-David et al., 2006; Rakhlin & Caponnetto, 2007; Shamir & Tishby, 2007; 2008) which investigates the convergence of optimal parameters both asymptotically and for finite samples. The vector quantization literature offers insights into the convergence of empirically optimal quantizers for k -Means in terms of the objective: A minimax rate of $\mathcal{O}\left(m^{-\frac{1}{2}}\right)$ is known if the underlying distribution has bounded support (Linder et al., 1994; Bartlett et al., 1998). A better rate of $\mathcal{O}\left(m^{-1}\right)$ may be achieved for finite support (Antos et al., 2005) or under both bounded support and regularity assumptions (Levrard et al., 2013). Ben-David (2007) provides a uniform convergence result for center based clustering under a bounded support assumption and Telgarsky & Dasgupta (2013) prove uniform deviation bounds for k -Means clustering if the underlying distribution satisfies moment assumptions (see Section 4).

Empirical risk minimization with fat-tailed losses has been studied in Mendelson et al. (2014), Mendelson (2014), Grünwald & Mehta (2016) and Dinh et al. (2016). Dinh et al. (2016) provide fast learning rates for k -Means clustering but under stronger assumptions than the ones considered in this paper. Guarantees similar to uniform deviation bounds can be obtained using importance sampling in the context of coresets construction (Bachem et al., 2015; Lucic et al., 2016; 2017).

Our contributions. We provide a new framework to obtain uniform deviation bounds for unbounded loss functions. We apply it to k -Means clustering and provide uniform deviation bounds with a rate of $\mathcal{O}\left(m^{-\frac{1}{2}}\right)$ for finite samples

under weak assumptions. In contrast to prior work, our bounds are all scale-invariant and hold for any set of k cluster centers (not *only* for a restricted solution set). We show that convergence depends on the kurtosis of the underlying distribution, which is the normalized fourth moment and measures the “tailedness” of a distribution. If bounded higher moments are available, we provide improved bounds that depend upon the normalized higher moments and we sharpen them even further under the stronger assumptions of subgaussianity and bounded support.

2. Problem statement for k -Means

We first focus on uniform deviation bounds for k -Means clustering, and defer the (more technical) framework for unbounded loss functions to Section 5. We consider the d -dimensional Euclidean space and define

$$d(x, Q)^2 = \min_{q \in Q} \|x - q\|_2^2$$

for any $x \in \mathbb{R}^d$ and any finite set $Q \subset \mathbb{R}^d$. Furthermore, slightly abusing the notation, for any $x, y \in \mathbb{R}^d$, we set $d(x, y)^2 = d(x, \{y\})^2 = \|x - y\|_2^2$.

Statistical k -Means. Let P be any distribution on \mathbb{R}^d with $\mu = \mathbb{E}_P[x]$ and $\sigma^2 = \mathbb{E}_P[d(x, \mu)^2] \in (0, \infty)$. For any set $Q \subset \mathbb{R}^d$ of $k \in \mathbb{N}$ cluster centers, the *expected quantization error* is given by $\mathbb{E}_P[d(x, Q)^2]$. The goal of the *statistical k -Means problem* is to find a set of k cluster centers such that the expected quantization error is minimized.

Empirical k -Means. Let \mathcal{X} denote a finite set of points in \mathbb{R}^d . The goal of the *empirical k -Means problem* is to find a set Q of k cluster centers in \mathbb{R}^d such that the *empirical quantization error* $\phi_{\mathcal{X}}(Q)$ is minimized, where

$$\phi_{\mathcal{X}}(Q) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} d(x, Q)^2.$$

Empirical risk minimization. In practice, the underlying data distribution P in statistical learning problems is often unknown. Instead, one is only able to observe independent samples from P . The *empirical risk minimization* principle advocates that finding a good solution on a random sample \mathcal{X}_m also provides provably good solutions to certain statistical learning problems if enough samples m are used.

Uniform deviation bounds. For k -Means, such a result may be shown by bounding the deviation between the expected loss and the empirical error, i.e.,

$$|\phi_{\mathcal{X}_m}(Q) - \mathbb{E}_P[d(x, Q)^2]|,$$

uniformly for all possible clusterings $Q \in \mathbb{R}^{d \times k}$. If this difference is sufficiently small for a given m , one may then solve the empirical k -Means problem on \mathcal{X}_m and obtain provable guarantees on the expected quantization error.

3. Uniform deviation bounds for k -Means

A simple approach would be to bound the deviation by an *absolute* error ϵ , i.e., to require that

$$|\phi_{\mathcal{X}_m}(Q) - \mathbb{E}_P[d(x, Q)^2]| \leq \epsilon \quad (1)$$

uniformly for a set of possible solutions (Telgarsky & Dasgupta, 2013). However, in this paper, we provide uniform deviation bounds of a more general form: For any distribution P and a sample of $m = f(\epsilon, \delta, k, d, P)$ points, we require that with probability at least $1 - \delta$

$$|\phi_{\mathcal{X}_m}(Q) - \mathbb{E}_P[d(x, Q)^2]| \leq \frac{\epsilon}{2}\sigma^2 + \frac{\epsilon}{2}\mathbb{E}_P[d(x, Q)^2] \quad (2)$$

uniformly for all $Q \in \mathbb{R}^{d \times k}$. The terms on the right-hand side may be interpreted as follows: The first term based on the variance σ^2 corresponds to a *scale-invariant, additive approximation error*. The second term is a *multiplicative approximation error* that allows the guarantee to hold even for solutions Q with a large expected quantization error. Similar additive error terms were recently explored by Bachem et al. (2016; 2017) in the context of seeding for k -Means.

There are three key reasons why we choose (2) over (1): First, (1) is not scale-invariant and may thus not hold for classes of distributions that are equal up to scaling. Second, (1) may not hold for an unbounded solution space, e.g. $\mathbb{R}^{d \times k}$. Third, we can always rescale P to unit variance and restrict ourselves to solutions Q with $\mathbb{E}_P[d(x, Q)^2] \leq \sigma^2$. Then, (2) implies (1) for a suitable transformation of P .

Importance of scale-invariance. If we scale all the points in a data set \mathcal{X} and all possible sets of solutions Q by some $\lambda > 0$, then the empirical quantization error is scaled by λ^2 . Similarly, if we consider the random variable λx where $x \sim P$, then the expected quantization error is scaled by λ^2 . At the same time, the k -Means problem remains the same: an optimal solution of the scaled problem is simply a scaled optimal solution of the original problem. Crucially, however, it is impossible to achieve the guarantee in (1) for distributions that are equal up to scaling: Suppose that (1) holds for some error tolerance ϵ , and sample size m with probability at least $1 - \delta$. Consider a distribution P and a solution $Q \in \mathbb{R}^{d \times k}$ such that with probability at least δ

$$a < |\phi_{\mathcal{X}_m}(Q) - \mathbb{E}_P[d(x, Q)^2]|.$$

for some $a > 0$.¹ For $\lambda > \frac{1}{\sqrt{a\epsilon}}$, let \tilde{P} be the distribution of the random variable λx where $x \sim P$ and let $\tilde{\mathcal{X}}_m$ consist of m samples from \tilde{P} . Defining $\tilde{Q} = \{\lambda q \mid q \in Q\}$, we have with probability at least δ

$$|\phi_{\tilde{\mathcal{X}}_m}(\tilde{Q}) - \mathbb{E}_{\tilde{P}}[d(x, \tilde{Q})^2]| > a\lambda^2 > \epsilon$$

¹For example, let P be a nondegenerate multivariate normal distribution and Q consist of k copies of the origin.

which contradicts (1) for the distribution \tilde{P} and the solution \tilde{Q} . Hence, (1) cannot hold for both P and its scaled transformation \tilde{P} .

Unrestricted solution space One way to guarantee *scale-invariance* would be require that

$$|\phi_{\mathcal{X}_m}(Q) - \mathbb{E}_P [d(x, Q)^2]| \leq \epsilon \sigma^2 \quad (3)$$

for all $Q \in \mathbb{R}^{d \times k}$. However, while (3) is scale-invariant, it is also impossible to achieve for *all* solutions Q as the following example shows. For simplicity, consider the 1-Means problem in 1-dimensional space and let P be a distribution with zero mean. Let \mathcal{X}_m denote m independent samples from P and denote by $\hat{\mu}$ the mean of \mathcal{X}_m . For any finite m , suppose that $\hat{\mu} \neq 0$ with high probability² and consider a solution Q consisting of a single point $q \in \mathbb{R}$. We then have

$$\begin{aligned} & |\phi_{\mathcal{X}_m}(\{q\}) - \mathbb{E}_P [d(x, \{q\})^2]| \\ &= |\phi_{\mathcal{X}_m}(\{\hat{\mu}\}) + d(\hat{\mu}, q)^2 - \sigma^2 - d(0, q)^2| \quad (4) \\ &= |\phi_{\mathcal{X}_m}(\{\hat{\mu}\}) - \sigma^2 + q^2 - 2q\hat{\mu} + \hat{\mu}^2 - q^2| \\ &= |\phi_{\mathcal{X}_m}(\{\hat{\mu}\}) - \sigma^2 + \hat{\mu}^2 - 2q\hat{\mu}| \end{aligned}$$

Since $\hat{\mu} \neq 0$ with high probability, clearly this expression diverges as $q \rightarrow \infty$ and thus (3) cannot hold for arbitrary solutions $Q \in \mathbb{R}^{d \times k}$. Intuitively, the key issue is that both the empirical and the statistical error become unbounded as $q \rightarrow \infty$. Previous approaches such as [Telgarsky & Dasgupta \(2013\)](#) solve this issue by restricting the solution space from $\mathbb{R}^{d \times k}$ to solutions that are no worse than some threshold. In contrast, we allow the deviation between the empirical and the expected quantization error to scale with $\mathbb{E}_P [d(x, Q)^2]$.

Arbitrary distributions. Finally, we show that we either need to impose assumptions on P or equivalently make the relationship between m , ϵ and δ in (2) depend on the underlying distribution P . Suppose that there exists a sample size $m \in \mathbb{N}$, an error tolerance $\epsilon \in (0, 1)$ and a maximal failure probability $\delta \in (0, 1)$ such that (2) holds for any distribution P . Let P be the Bernoulli distribution on $\{0, 1\} \subset \mathbb{R}$ with $\mathbb{P}[x = 1] = p$ for $p \in (\delta^{\frac{1}{m}}, 1)$. By design, we have $\mu = p$, $\sigma^2 = p(1-p)$ and $\mathbb{E}_P [d(x, 1)^2] = (1-p)$. Furthermore, with probability at least δ , the set \mathcal{X}_m of m independent samples from P consists of m copies of a point at one. Hence, (2) implies that with probability at least $1 - \delta$

$$|\phi_{\mathcal{X}_m}(1) - \mathbb{E}_P [d(x, 1)^2]| \leq \epsilon \mathbb{E}_P [d(x, 1)^2]$$

since $\sigma^2 \leq \mathbb{E}_P [d(x, 1)^2]$. However, with probability at least δ , we have $\phi_{\mathcal{X}_m}(1) = 0$ which would imply $1 \leq \epsilon$ and thus lead to a contradiction with $\epsilon \in (0, 1)$.

²For example, if P is the standard normal distribution.

4. Key results for k -Means

In this section, we present our main results for k -Means and defer the analysis and proofs to Sections 6.

4.1. Kurtosis bound

Similar to [Telgarsky & Dasgupta \(2013\)](#), the weakest assumption that we require is that the fourth moment of $d(x, \mu)$ for $x \in P$ is bounded.³ Our results are based on the *kurtosis* of P which we define as

$$\hat{M}_4 = \frac{\mathbb{E}_P [d(x, \mu)^4]}{\sigma^4}.$$

The kurtosis is the normalized fourth moment and is a scale-invariant measure of the ‘‘tailedness’’ of a distribution. For example, the normal distribution has a kurtosis of 3, while more heavy tailed distributions such as the t -Student distribution or the Pareto distribution have a potentially unbounded kurtosis. A natural interpretation of the kurtosis is provided by [Moors \(1986\)](#). For simplicity, consider a data set with unit variance. Then, the kurtosis may be restated as the shifted variance of $d(x, \mu)^2$, i.e.,

$$\hat{M}_4 = \text{Var} (d(x, \mu)^2) + 1.$$

This provides a valuable insight into why the kurtosis is relevant for our setting: For simplicity, suppose we would like to estimate the expected quantization error $\mathbb{E}_P [d(x, \mu)^2]$ by the empirical quantization error $\phi_{\mathcal{X}_m}(\{\mu\})$ on a finite sample \mathcal{X}_m .⁴ Then, the kurtosis measures the dispersion of $d(x, \mu)^2$ around its mean $\mathbb{E}_P [d(x, \mu)^2]$ and provides a bound on how many samples are required to achieve an error of ϵ . While this simple example provides the key insight for the trivial solution $Q = \{\mu\}$, it requires a non-trivial effort to extend the guarantee in (2) to hold uniformly for all solutions $Q \in \mathbb{R}^{d \times k}$.

With the use of a novel framework to learn unbounded loss functions (presented in Section 5), we are able to provide the following guarantee for k -Means.

Theorem 1 (Kurtosis). *Let $\epsilon, \delta \in (0, 1)$ and $k \in \mathbb{N}$. Let P be any distribution on \mathbb{R}^d with kurtosis $\hat{M}_4 < \infty$. For*

$$m \geq \frac{12800 \left(8 + \hat{M}_4\right)}{\epsilon^2 \delta} \left(3 + 30k(d + 4) \log 6k + \log \frac{1}{\delta}\right)$$

let $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$ be m independent samples from P . Then, with probability at least $1 - \delta$, for all $Q \in \mathbb{R}^{d \times k}$

$$|\phi_{\mathcal{X}}(Q) - \mathbb{E}_P [d(x, Q)^2]| \leq \frac{\epsilon}{2} \sigma^2 + \frac{\epsilon}{2} \mathbb{E}_P [d(x, Q)^2].$$

³While our random variables $x \in P$ are potentially multivariate, it suffices to consider the behavior of the univariate random variable $d(x, \mu)$ for the assumptions in this section.

⁴This is a hypothetical exercise as $\mathbb{E}_P [d(x, \mu)^2] = 1$ by design. However, it provides an insight to the importance of the kurtosis.

The proof is provided in Section 6.1. The number of sufficient samples

$$m \in \mathcal{O} \left(\frac{\hat{M}_4}{\epsilon^2 \delta} \left(dk \log k + \log \frac{1}{\delta} \right) \right)$$

is linear in the kurtosis \hat{M}_4 and the dimensionality d , near-linear in the number of clusters k and $\frac{1}{\delta}$, and quadratic in $\frac{1}{\epsilon}$. Intuitively, the bound may be interpreted as follows: $\Omega \left(\frac{\hat{M}_4}{\epsilon^2 \delta} \right)$ samples are required such that the guarantee holds for a single solution $Q \in \mathbb{R}^{d \times k}$. Informally, a generalization of the Vapnik Chervonenkis dimension for k -Means clustering may be bounded by $\mathcal{O}(dk \log k)$ and measures the ‘‘complexity’’ of the learning problem. The multiplicative $dk \log k + \log \frac{1}{\delta}$ term intuitively extends the guarantee uniformly to all possible $Q \in \mathbb{R}^{d \times k}$.

Comparison to Telgarsky & Dasgupta (2013). While we require a bound on the normalized fourth moment, i.e., the kurtosis, Telgarsky & Dasgupta (2013) consider the case where all unnormalized moments up to the fourth are uniformly bounded by some M , i.e.,

$$\mathbb{E}_P [d(x, \mu)^l] \leq M, \quad 1 \leq l \leq 4.$$

They provide uniform deviation bounds for all solutions Q such that either $\phi_{\mathcal{X}}(Q) \leq c$ or $\mathbb{E}_P [d(x, Q)^2] \leq c$ for some $c > 0$. To compare our bounds, we consider a data set with unit variance and restrict ourselves to solutions $Q \in \mathbb{R}^{d \times k}$ with an expected quantization error of at most the variance, i.e., $\mathbb{E}_P [d(x, Q)^2] \leq \sigma^2 = 1$. Consider the deviation

$$\Delta = \sup_{Q \in \mathbb{R}^{d \times k}: \mathbb{E}_P [d(x, Q)^2] \leq 1} |\phi_{\mathcal{X}}(Q) - \mathbb{E}_P [d(x, Q)^2]|.$$

Telgarsky & Dasgupta (2013) bound this deviation by

$$\Delta \in \mathcal{O} \left(\sqrt{\frac{M^2}{\sqrt{m}} \left(dk \log(Mdm) + \log \frac{1}{\delta} \right)} + \sqrt{\frac{1}{m\delta^2}} \right).$$

In contrast, our bound in Theorem 1 implies

$$\Delta \in \mathcal{O} \left(\sqrt{\frac{\hat{M}_4}{m\delta} \left(dk \log k + \log \frac{1}{\delta} \right)} \right).$$

The key difference lies in how Δ scales with the sample size m . While Telgarsky & Dasgupta (2013) show a rate of $\Delta \in \mathcal{O}(m^{-\frac{1}{4}})$, we improve it to $\Delta \in \mathcal{O}(m^{-\frac{1}{2}})$.

4.2. Bounded higher moments

The tail behavior of $d(x, \mu)$ may be characterized by the moments of P . For $p \in \mathbb{N}$, consider the *standardized p -th moment* of P , i.e.,

$$\hat{M}_p = \frac{\mathbb{E}_P [d(x, \mu)^p]}{\sigma^p}.$$

Theorem 2 provides an improved uniform deviation bound if P has bounded higher moments.

Theorem 2 (Moment bound). *Let $\epsilon \in (0, 1)$, $\delta \in (0, 1)$ and $k \in \mathbb{N}$. Let P be any distribution on \mathbb{R}^d with finite p -th order moment bound $\hat{M}_p < \infty$ for $p \in \{4, 8, \dots, \infty\}$. For $m \geq \max \left(\frac{3200m_1}{\epsilon^2}, \left(\frac{8}{\delta} \right)^{\frac{8}{p}} \right)$ with*

$$m_1 = p \left(4 + \hat{M}_p^{\frac{4}{p}} \right) \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta} \right)$$

let $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$ be m independent samples from P . Then, with probability at least $1 - \delta$, for all $Q \in \mathbb{R}^{d \times k}$

$$|\phi_{\mathcal{X}}(Q) - \mathbb{E}_P [d(x, Q)^2]| \leq \frac{\epsilon}{2} \sigma^2 + \frac{\epsilon}{2} \mathbb{E}_P [d(x, Q)^2].$$

The proof is provided in Section 6.2. Compared to the previous bound based on the kurtosis, Theorem 2 requires

$$m \in \mathcal{O} \left(\frac{p \hat{M}_p^{\frac{4}{p}}}{\epsilon^2} \left(dk \log k + \log \frac{1}{\delta} \right) + \left(\frac{1}{\delta} \right)^{\frac{8}{p}} \right)$$

samples. With higher order moment bounds, it is easier to achieve high probability results since the dependence on $\frac{1}{\delta}$ is only of $\mathcal{O} \left(\left(\frac{1}{\delta} \right)^{\frac{8}{p}} \right)$ compared to near linear for a kurtosis

bound. The quantity $\hat{M}_p^{\frac{4}{p}}$ may be interpreted as a bound on the kurtosis \hat{M}_4 based on the higher order moment \hat{M}_p since Hoelder’s inequality implies that $\hat{M}_4 \leq \hat{M}_p^{\frac{4}{p}}$. While the result only holds for $p \in \{8, 12, 16, \dots, \infty\}$, it is trivially extended to $p' \geq 8$: Consider Theorem 2 with $p = 4 \left\lfloor \frac{p'}{4} \right\rfloor$ and note that by Hoelder’s inequality $\hat{M}_p^{\frac{4}{p}} \leq \hat{M}_{p'}^{\frac{4}{p'}}$.

Comparison to Telgarsky & Dasgupta (2013). Again, we consider distributions P that have unit variance and we restrict ourselves to solutions $Q \in \mathbb{R}^{d \times k}$ with an expected quantization error of at most the variance, i.e., $\mathbb{E}_P [d(x, Q)^2] \leq \sigma^2 = 1$. Telgarsky & Dasgupta (2013) require that there exists a bound M

$$\mathbb{E}_P [d(x, \mu)^l] \leq M, \quad 1 \leq l \leq p.$$

Then, for m sufficiently large, Δ is of

$$\mathcal{O} \left(\sqrt{\frac{M^{\frac{8}{p}}}{m^{1-\frac{4}{p}}} \left(dk \ln(M^{\frac{4}{p}} dm) + \ln \frac{1}{\delta} \right)} + \frac{2^{\frac{p}{4}}}{m^{\frac{3}{4}-\frac{2}{p}}} \left(\frac{1}{\delta} \right)^{\frac{4}{p}} \right).$$

In contrast, we obtain that, for m sufficiently large,

$$\Delta \in \mathcal{O} \left(\sqrt{\frac{p \hat{M}_p^{\frac{4}{p}}}{m} \left(dk \log k + \log \frac{1}{\delta} \right)} \right).$$

While Telgarsky & Dasgupta (2013) only show a rate of $\mathcal{O}(m^{-\frac{1}{2}})$ as $p \rightarrow \infty$, we obtain a $\mathcal{O}(m^{-\frac{1}{2}})$ rate for all higher moment bounds.

4.3. Subgaussianity

If the distribution P is subgaussian, then all its moments \hat{M}_p are bounded. By optimizing p in Theorem 2, we are able to show the following bound.

Theorem 3 (Subgaussian bound). *Let $\epsilon \in (0, 1)$, $\delta \in (0, 1)$ and $k \in \mathbb{N}$. Let P be any distribution on \mathbb{R}^d with $\mu = \mathbb{E}_P[x]$ and*

$$\forall t > 0 : \quad \mathbb{P}[d(x, \mu) > t\sigma] \leq a \exp\left(-\frac{t^2}{\sqrt{b}}\right)$$

for some $a > 1, b > 0$. Let $m \geq \frac{3200m_1}{\epsilon^2}$ with

$$m_1 = p \left(4 + \frac{abp^2}{4}\right) \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta}\right).$$

and $p = 9 + 3 \log \frac{1}{\delta}$. Let $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$ be m independent samples from P . Then, with probability at least $1 - \delta$, for all $Q \in \mathbb{R}^{d \times k}$

$$|\phi_{\mathcal{X}}(Q) - \mathbb{E}_P[d(x, Q)^2]| \leq \frac{\epsilon}{2}\sigma^2 + \frac{\epsilon}{2}\mathbb{E}_P[d(x, Q)^2].$$

The proof is provided in Section 6.3. In $\mathcal{O}(\cdot)$ notation,

$$m \in \mathcal{O}\left(\frac{ab \log^3 \frac{1}{\delta}}{\sigma \epsilon^2} \left(dk \log k + \log \frac{1}{\delta}\right)\right)$$

samples are hence sufficient. This result features a polylogarithmic dependence on $\frac{1}{\delta}$ compared to the polynomial dependence for the bounds based on bounded higher moments. The sufficient sample size further scales linearly with the (scale-invariant) subgaussianity parameters a and b . For example, if P is a one-dimensional normal distribution of any scale, we would have $a = 2$ and $b = 1$.

4.4. Bounded support

The strongest assumption that we consider is if the support of P is bounded by a hypersphere in \mathbb{R}^d with diameter $R > 0$. This ensures that almost surely $d(x, \mu) \leq R$ and hence $\hat{M}_4 \leq \frac{R^4}{\sigma^4}$. This allows us to obtain Theorem 4.

Theorem 4 (Bounded support). *Let $\epsilon \in (0, 1)$, $\delta \in (0, 1)$ and $k \in \mathbb{N}$. Let P be any distribution on \mathbb{R}^d , with $\mu = \mathbb{E}_P[x]$ and $\sigma^2 = \mathbb{E}_P[d(x, \mu)^2] \in (0, \infty)$, whose support is contained in a d -dimensional hypersphere of diameter $R > 0$. For*

$$m \geq \frac{12800 \left(8 + \frac{R^4}{\sigma^4}\right)}{\epsilon^2} \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta}\right)$$

let $\mathcal{X} = \{x_1, x_2, \dots, x_m\}$ be m independent samples from P . Then, with probability at least $1 - \delta$, for all $Q \in \mathbb{R}^{d \times k}$

$$|\phi_{\mathcal{X}}(Q) - \mathbb{E}_P[d(x, Q)^2]| \leq \frac{\epsilon}{2}\sigma^2 + \frac{\epsilon}{2}\mathbb{E}_P[d(x, Q)^2].$$

The proof is provided in Section 6.4. Again, the sufficient sample size scales linearly with the kurtosis bound $\frac{R^4}{\sigma^4}$. However, the bound is only logarithmic in $\frac{1}{\delta}$.

5. Framework for unbounded loss functions

To obtain the results presented in Section 4, we propose a novel framework to uniformly approximate the expected values of a set of unbounded functions based on a random sample. We consider a function family \mathcal{F} mapping from an arbitrary input space \mathcal{X} to $\mathbb{R}_{\geq 0}$ and a distribution P on \mathcal{X} . We further require a generalization of the *Vapnik-Chervonenkis dimension* to continuous, unbounded functions⁵ — the *pseudo-dimension*.

Definition 1 (Haussler (1992); Li et al. (2001)). *The pseudo-dimension of a set \mathcal{F} of functions from \mathcal{X} to $\mathbb{R}_{\geq 0}$, denoted by $\text{Pdim}(\mathcal{F})$, is the largest d' such there is a sequence $x_1, \dots, x_{d'}$ of domain elements from \mathcal{X} and a sequence $r_1, \dots, r_{d'}$ of reals such that for each $b_1, \dots, b_{d'} \in \{\text{above}, \text{below}\}$, there is an $f \in \mathcal{F}$ such that for all $i = 1, \dots, d'$, we have $f(x_i) \geq r_i \iff b_i = \text{above}$.*

Similar to the VC dimension, the pseudo-dimension measures the cardinality of the largest subset of \mathcal{X} that can be *shattered* by the function family \mathcal{F} . Informally, the pseudo-dimension measures the richness of \mathcal{F} and plays a critical role in providing a uniform approximation guarantee across all $f \in \mathcal{F}$. With this notion, we are able to state the main result in our framework.

Theorem 5. *Let $\epsilon \in (0, 1)$, $\delta \in (0, 1)$ and $t > 0$. Let \mathcal{F} be a family of functions from \mathcal{X} to $\mathbb{R}_{\geq 0}$ with $\text{Pdim}(\mathcal{F}) = d < \infty$. Let $s : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ be a function such that $s(x) \geq \sup_{f \in \mathcal{F}} f(x)$ for all $x \in \mathcal{X}$. Let P be any distribution on \mathcal{X} and for*

$$m \geq \frac{200t}{\epsilon^2} \left(3 + 5d + \log \frac{1}{\delta}\right),$$

let x_1, x_2, \dots, x_{2m} be $2m$ independent samples from P .

Then, if

$$\mathbb{E}_P[s(x)^2] \leq t \quad \text{and} \quad \mathbb{P}\left[\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 > t\right] \leq \frac{\delta}{4}, \quad (5)$$

it holds with probability at least $1 - \delta$ that

$$\left|\frac{1}{m} \sum_{i=1}^m f(x_i) - \mathbb{E}_P[f(x)]\right| \leq \epsilon, \quad \forall f \in \mathcal{F}. \quad (6)$$

Applying Theorem 5 to a function family \mathcal{F} requires three steps: First, one needs to bound the pseudo-dimension of

⁵The pseudo-dimension was originally defined for sets of functions mapping to $[0, 1]$ (Haussler, 1992; Li et al., 2001). However, it is trivially extended to unbounded functions mapping to $\mathbb{R}_{\geq 0}$.

\mathcal{F} . Second, it is necessary to find a function $s : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ such that

$$f(x) \leq s(x), \quad \forall x \in \mathcal{X} \text{ and } \forall f \in \mathcal{F}.$$

Ideally, such a bound should be as tight as possible. Third, one needs to find some $t > 0$ and a sample size

$$m \geq \frac{200t}{\epsilon^2} \left(3 + 5d + \log \frac{1}{\delta} \right)$$

such that

$$\mathbb{E}_P \left[s(x)^2 \right] \leq t \quad \text{and} \quad \mathbb{P} \left[\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 > t \right] \leq \frac{\delta}{4}.$$

Finding such a bound usually entails examining the tail behavior of $s(x)^2$ under P . Furthermore, it is evident that a bound t may only be found if $\mathbb{E}_P \left[s(x)^2 \right]$ is bounded and that assumptions on the distribution P are required. In Section 6, we will see that for k -Means a function $s(x)$ with $\mathbb{E}_P \left[s(x)^2 \right] < \infty$ may be found if the kurtosis of P is bounded.

We defer the proof of Theorem 5 to Section B of the Supplementary Materials and provide a short proof sketch that captures the main insight. The proof is based on *symmetrization*, the bounding of *covering numbers* and *chaining* — common techniques in the empirical process literature (Pollard, 1984; Li et al., 2001; Boucheron et al., 2013; Koltchinskii, 2011; van der Vaart & Wellner, 1996). The novelty lies in considering loss functions $f(\cdot)$ and cover functions $s(\cdot)$ in Theorem 5 that are potentially unbounded.

Proof sketch. Our proof is based on a *double sampling* approach. Let $x_{m+1}, x_{m+2}, \dots, x_{2m}$ be an additional m independent samples from P and let $\sigma_1, \sigma_2, \dots, \sigma_m$ be independent random variables uniformly sampled from $\{-1, 1\}$. Then, we show that, if $\mathbb{E}_P \left[s(x)^2 \right] \leq t$, the probability of (6) not holding may be bounded by the probability that there exists a $f \in \mathcal{F}$ such that

$$\left| \frac{1}{m} \sum_{i=1}^m \sigma_i (f(x_i) - f(x_{i+m})) \right| > \epsilon. \quad (7)$$

We first provide the intuition for a single function $f \in \mathcal{F}$ and then show how we extend it to all $f \in \mathcal{F}$. While the function $f(x)$ is not bounded, for a given sample x_1, x_2, \dots, x_{2m} , each $f(x_i)$ is contained within $[0, s(x_i)]$. Given the sample x_1, x_2, \dots, x_{2m} , the random variable $\sigma_i (f(x_i) - f(x_{i+m}))$ is bounded in $0 \pm \max(s(x_i), s(x_{i+m}))$ and has zero mean. Hence, given independent samples x_1, x_2, \dots, x_{2m} , the probability of (7) occurring for a single $f \in \mathcal{F}$ can be

bounded using Hoeffding's inequality by

$$2 \exp \left(- \frac{2m\epsilon^2}{\frac{1}{m} \sum_{i=1}^m \max(s(x_i), s(x_{i+m}))^2} \right) \leq 2 \exp \left(- \frac{m\epsilon^2}{\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2} \right).$$

By (5), with probability at least $1 - \frac{\delta}{4}$, we have $\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 \leq t$ and we hence require $m \in \Omega \left(\frac{t \log \frac{1}{\delta}}{\epsilon^2} \right)$ samples to guarantee that (7) does not hold for a single $f \in \mathcal{F}$ with probability at least $1 - \frac{\delta}{2}$.

To bound the probability that there exists *any* $f \in \mathcal{F}$ such that (7) holds, we show in Lemma 5 (see Section B of the Supplementary Materials) that, given independent samples x_1, x_2, \dots, x_{2m} ,

$$\mathbb{P} \left[\exists f \in \mathcal{F} : \left| \frac{1}{m} \sum_{i=1}^m \sigma_i (f(x_i) - f(x_{i+m})) \right| > \epsilon \right] \leq 4 (16e^2)^{\text{Pdim}(\mathcal{F})} e^{-\frac{\epsilon^2 m}{200 \frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2}}.$$

The key difficulty in proving Lemma 5 is that the functions $f \in \mathcal{F}$ are not bounded uniformly in $[0, 1]$. To this end, we provide in Lemma 4 a novel result that bounds the size of ϵ -packings of \mathcal{F} if the functions $f \in \mathcal{F}$ are bounded in expectation. Based on Lemma 5, we then prove the main claim of Theorem 5. \square

6. Analysis for k -Means

In order to apply Theorem 5 to k -Means clustering, we require a suitable family \mathcal{F} , an upper bound $s(x)$ and a bound on $\mathbb{E}_P \left[s(x)^2 \right]$. We provide this in Lemma 1 and defer bounding $\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2$ to the proofs of Theorems 2-4.

Lemma 1 (k -Means). *Let $k \in \mathbb{N}$. Let P be any distribution on \mathbb{R}^d with $\mu = \mathbb{E}_P[x]$, $\sigma^2 = \mathbb{E}_P[d(x, \mu)^2] \in (0, \infty)$ and bounded kurtosis \hat{M}_4 . For any $x \in \mathbb{R}^d$ and any $Q \in \mathbb{R}^{d \times k}$, define*

$$f_Q(x) = \frac{d(x, Q)^2}{\frac{1}{2}\sigma^2 + \frac{1}{2}\mathbb{E}_P[d(x, Q)^2]} \quad (8)$$

as well as the function family $\mathcal{F} = \{f_Q(\cdot) \mid Q \in \mathbb{R}^{d \times k}\}$. Let

$$s(x) = \frac{4d(x, \mu)^2}{\sigma^2} + 8.$$

We then have

$$\text{Pdim}(\mathcal{F}) \leq 6k(d+4) \log 6k, \quad (9)$$

$$f_Q(x) \leq s(x) \quad (10)$$

for any $x \in \mathbb{R}^d$ and $Q \in \mathbb{R}^{d \times k}$ and

$$\mathbb{E}_P [s(x)^2] = 128 + 16\hat{M}_4. \quad (11)$$

The proof of Lemma 1 is provided in Section C of the Supplementary Materials. The definition of $f_Q(x)$ in (8) is motivated as follows: If we use Theorem 5 to guarantee

$$\left| \sum_{i=1}^m f(x_i) - \mathbb{E}_P [f(x)] \right| \leq \epsilon \quad \forall f \in \mathcal{F}. \quad (12)$$

then this implies

$$|\phi_{\mathcal{X}}(Q) - \mathbb{E}_P [d(x, Q)^2]| \leq \frac{\epsilon}{2}\sigma^2 + \frac{\epsilon}{2}\mathbb{E}_P [d(x, Q)^2] \quad (13)$$

as is required by Theorems 2-4. Lemma 1 further shows that $\mathbb{E} [s(x)^2]$ is bounded if and only if the kurtosis of P is bounded. This is the reason why a bounded kurtosis is the weakest assumption on P that we require in Section 4.

We now proceed to prove Theorems 2-4 by applying Theorem 5 and examining the tail behavior of $\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2$.

6.1. Proof of Theorem 1 (kurtosis bound)

The bound based on the kurtosis follows easily from Markov's inequality.

Proof. We consider the choice $t = 4(128 + 16\hat{M}_4)/\delta$. By Markov's inequality and linearity of expectation, we then have by Lemma 1 that

$$\mathbb{P} \left[\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 > t \right] \leq \frac{\mathbb{E} [s(x)^2]}{t} = \frac{\delta}{4}.$$

Furthermore, $\mathbb{E}_P [s(x)^2] \leq t$. Hence, we may apply Theorem 5 to obtain that for

$$m \geq \frac{12800(8 + \hat{M}_4)}{\epsilon^2 \delta} \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta} \right),$$

it holds with probability at least $1 - \delta$ that

$$\left| \frac{1}{m} \sum_{i=1}^m f(x_i) - \mathbb{E} [f(x)] \right| \leq \epsilon \quad \forall f \in \mathcal{F}.$$

This implies the main claim and thus concludes the proof. \square

6.2. Proof of Theorem 2 (higher order moment bound)

We prove the result by bounding the higher moments of $\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2$ using the Marcinkiewicz-Zygmund inequality and subsequently applying Markov's inequality.

Proof. Hoelder's inequality implies

$$\hat{M}_4 = \frac{\mathbb{E}_P [d(x, \mu)^4]}{\sigma^4} \leq \frac{\mathbb{E}_P [d(x, \mu)^p]^{\frac{4}{p}}}{\sigma^4} \leq \hat{M}_p^{\frac{4}{p}}$$

Hence, by Lemma 1 we have that $\mathbb{E}_P [s(x)^2] \leq 128 + 16\hat{M}_p^{\frac{4}{p}}$. Since $s(x)^2 \geq 0$ for all $x \in \mathbb{R}^d$, we have

$$\begin{aligned} |s(x)^2 - \mathbb{E}_P [s(x)^2]| &\leq \max(s(x)^2, \mathbb{E}_P [s(x)^2]) \\ &\leq \max\left(s(x)^2, 128 + 16\hat{M}_p^{\frac{4}{p}}\right) \\ &\leq 128 \\ &\quad + 16 \max\left(\hat{M}_p^{\frac{4}{p}}, 2\frac{d(x, \mu)^4}{\sigma^4}\right). \end{aligned} \quad (14)$$

This implies that

$$\begin{aligned} \mathbb{E}_P \left[|s(x)^2 - \mathbb{E}_P [s(x)^2]|^{\frac{p}{4}} \right] &\leq 256^{\frac{p}{4}} + 32^{\frac{p}{4}} \max\left(\hat{M}_p, 2^{\frac{p}{4}} \frac{\mathbb{E}_P [d(x, \mu)^p]}{\sigma^p}\right) \\ &\leq 256^{\frac{p}{4}} + 32^{\frac{p}{4}} \max\left(\hat{M}_p, 2^{\frac{p}{4}} \hat{M}_p\right) \\ &\leq 256^{\frac{p}{4}} + 64^{\frac{p}{4}} \hat{M}_p. \end{aligned} \quad (15)$$

We apply a variant of the Marcinkiewicz-Zygmund inequality (Ren & Liang, 2001) to the zero-mean random variable $s(x)^2 - \mathbb{E}_P [s(x)^2]$ to obtain

$$\begin{aligned} \mathbb{E}_P \left[\left| \frac{1}{2m} \sum_{i=1}^{2m} (s(x_i)^2 - \mathbb{E}_P [s(x)^2]) \right|^{\frac{p}{4}} \right] &\leq \left(\frac{p-4}{4\sqrt{2m}} \right)^{\frac{p}{4}} \mathbb{E}_P \left[|s(x)^2 - \mathbb{E}_P [s(x)^2]|^{\frac{p}{4}} \right] \\ &\leq \left(\frac{p-4}{4\sqrt{2m}} \right)^{\frac{p}{4}} (256^{\frac{p}{4}} + 64^{\frac{p}{4}} \hat{M}_p) \end{aligned} \quad (16)$$

For $u > 0$, the Markov inequality implies

$$\begin{aligned} \mathbb{P} \left[\left| \frac{1}{2m} \sum_{i=1}^{2m} (s(x_i)^2 - \mathbb{E}_P [s(x)^2]) \right| > u \right] &\leq \left(\frac{p-4}{4u\sqrt{2m}} \right)^{\frac{p}{4}} (256^{\frac{p}{4}} + 64^{\frac{p}{4}} \hat{M}_p) \\ &\leq \left(\frac{p-4}{4u\sqrt{2m}} \right)^{\frac{p}{4}} 2 \max(256^{\frac{p}{4}}, 64^{\frac{p}{4}} \hat{M}_p) \\ &\leq 2 \left(\frac{p-4}{4u\sqrt{2m}} \max(256, 64\hat{M}_p^{\frac{4}{p}}) \right)^{\frac{p}{4}} \\ &\leq 2 \left(\frac{p-4}{u\sqrt{2m}} (64 + 16\hat{M}_p^{\frac{4}{p}}) \right)^{\frac{p}{4}}. \end{aligned} \quad (17)$$

For $u = (p-4) \left(64 + 16\hat{M}_p^{\frac{4}{p}}\right)$, we thus have

$$\mathbb{P} \left[\left| \frac{1}{2m} \sum_{i=1}^{2m} (s(x_i)^2 - \mathbb{E}_P [s(x)^2]) \right| > u \right] \leq 2m^{-\frac{p}{8}} \quad (18)$$

Since $m \geq \left(\frac{8}{\delta}\right)^{\frac{p}{8}}$, this implies

$$\mathbb{P} \left[\left| \frac{1}{2m} \sum_{i=1}^{2m} (s(x_i)^2 - \mathbb{E}_P [s(x)^2]) \right| > u \right] \leq \frac{\delta}{4} \quad (19)$$

It holds that

$$\begin{aligned} u + \mathbb{E}_P [s(x)^2] &= (p-4) \left(64 + 16\hat{M}_p^{\frac{4}{p}}\right) + 128 + 16\hat{M}_p^{\frac{4}{p}} \\ &\leq p \left(64 + 16\hat{M}_p^{\frac{4}{p}}\right) \end{aligned} \quad (20)$$

We set $t = p \left(64 + 16\hat{M}_p^{\frac{4}{p}}\right)$ and thus have

$$\mathbb{P} \left[\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 > t \right] \leq \frac{\delta}{4} \quad (21)$$

In combination with $\mathbb{E}_P [s(x)^2] \leq t$ by Lemma 1, we may thus apply Theorem 5. Since $m \geq \frac{3200m_1}{\epsilon^2}$ with

$$m_1 = p \left(4 + \hat{M}_p^{\frac{4}{p}}\right) \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta}\right)$$

it holds with probability at least $1 - \delta$ that

$$\left| \frac{1}{m} \sum_{i=1}^m f(x_i) - \mathbb{E} [f(x)] \right| \leq \epsilon \quad \forall f \in \mathcal{F}.$$

This implies the main claim and thus concludes the proof. \square

6.3. Proof of Theorem 3 (subgaussianity)

Under subgaussianity, all moments of $d(x, \mu)$ are bounded. We show the result by optimizing over p in Theorem 2.

Proof. For $p \in \{4, 8, \dots, \infty\}$, we have

$$\begin{aligned} \hat{M}_p &= \mathbb{E}_P \left[\left| \frac{d(x, \mu)}{\sigma} \right|^p \right] \\ &= \int_0^\infty \mathbb{P} \left[\frac{d(x, \mu)}{\sigma} > u^{\frac{1}{p}} \right] du \\ &\leq \int_0^\infty a \exp \left(-\frac{u^{\frac{2}{p}}}{\sqrt{b}} \right) du. \end{aligned} \quad (22)$$

Let $u(t) = b^{\frac{p}{4}} t^{\frac{p}{2}}$ which implies $du/dt = b^{\frac{p}{4}} \frac{p}{2} t^{\frac{p}{2}-1}$. Hence,

$$\hat{M}_p \leq \frac{ab^{\frac{p}{4}}}{2} \int_0^\infty e^{-t} t^{\frac{p}{2}-1} dt.$$

By the definition of the gamma function and since p is even, we have

$$\int_0^\infty e^{-t} t^{\frac{p}{2}-1} dt = \Gamma \left(\frac{p}{2} \right) = \left(\frac{p}{2} - 1 \right)! \leq \left(\frac{p}{2} \right)^{\frac{p}{2}-1}$$

Hence, for $p \in \{4, 8, \dots, \infty\}$, we have

$$\hat{M}_p^{\frac{4}{p}} \leq \frac{1}{4} a^{\frac{4}{p}} b p^2 \leq \frac{1}{4} a b p^2.$$

Let $p_* = 4 \left\lceil \frac{5}{4} + \frac{3}{4} \log \frac{1}{\delta} \right\rceil$ which implies

$$p_* \geq 5 + 3 \log \frac{1}{\delta} \geq \frac{8}{\log 48} \log \frac{8}{\delta}$$

and thus $\left(\frac{8}{\delta}\right)^{\frac{p_*}{8}} \leq 48$. We instantiate Theorem 2 with the p_* th-order bound \hat{M}_{p_*} of P . Since $\left(\frac{8}{\delta}\right)^{\frac{p_*}{8}} \leq 48$, the minimum sample size is thus

$$\frac{3200p_*}{\epsilon^2} \left(4 + \frac{abp_*^2}{4}\right) \left(3 + 30k(d+4) \log 6k + \log \frac{1}{\delta}\right).$$

The main claim finally holds since $p_* \leq p = 9 + 3 \log \frac{1}{\delta}$. \square

6.4. Proof of Theorem 4 (bounded support)

Proof. Let $t = 128 + 64R^4/\sigma^4$. Since the support of P is bounded, we have $s(x) \leq t$ for all $x \in \mathbb{R}^d$. This implies that $\mathbb{E}_P [s(x)^2] \leq t$ and that $\frac{1}{2m} \sum_{i=1}^{2m} s(x_i)^2 \leq t$ almost surely. The result then follows from Theorem 5. \square

7. Conclusion

We have presented a framework to uniformly approximate the expected value of unbounded functions on a sample. With this framework we are able to provide theoretical guarantees for empirical risk minimization in k -Means clustering if the kurtosis of the underlying distribution is bounded. In particular, we obtain state-of-the-art bounds on the sufficient number of samples to achieve a given uniform approximation error. If the underlying distribution fulfills stronger assumptions, such as bounded higher moments, subgaussianity or bounded support, our analysis yields progressively better bounds. We conjecture that Theorem 5 can be applied to other related problems such as hard and soft Bregman clustering, likelihood estimation of Gaussian mixture models, as well as nonparametric clustering problems. However, such results do not follow immediately and require additional arguments beyond the scope of this paper.

Acknowledgements

This research was partially supported by SNSF NRP 75, ERC StG 307036, a Google Ph.D. Fellowship and an IBM Ph.D. Fellowship. This work was done in part while Andreas Krause was visiting the Simons Institute for the Theory of Computing.

References

- Antos, András, Györfi, L., and Gyögy, Andras. Individual convergence rates in empirical vector quantizer design. *IEEE Transactions on Information Theory*, 51(11):4013–4022, 2005.
- Bachem, Olivier, Lucic, Mario, and Krause, Andreas. Coresets for nonparametric estimation - the case of DP-means. In *International Conference on Machine Learning (ICML)*, 2015.
- Bachem, Olivier, Lucic, Mario, Hassani, S. Hamed, and Krause, Andreas. Fast and provably good seedings for k -means. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 55–63, 2016.
- Bachem, Olivier, Lucic, Mario, and Krause, Andreas. Distributed and provably good seedings for k -means in constant rounds. In *To appear in International Conference on Machine Learning (ICML)*, pp. 209–217, 2017.
- Bartlett, Peter L, Linder, Tamás, and Lugosi, Gábor. The minimax distortion redundancy in empirical quantizer design. *IEEE Transactions on Information Theory*, 44(5):1802–1813, 1998.
- Ben-David, Shai. A framework for statistical clustering with constant time approximation algorithms for k -median and k -means clustering. *Machine Learning*, 66(2-3):243–257, 2007.
- Ben-David, Shai, Von Luxburg, Ulrike, and Pál, Dávid. A sober look at clustering stability. In *International Conference on Computational Learning Theory (COLT)*, pp. 5–19. Springer, 2006.
- Boucheron, Stéphane, Lugosi, Gábor, and Massart, Pascal. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford University Press, 2013.
- Dinh, Vu C, Ho, Lam S, Nguyen, Binh, and Nguyen, Duy. Fast learning rates with heavy-tailed losses. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 505–513, 2016.
- Grünwald, Peter D and Mehta, Nishant A. Fast rates with unbounded losses. *arXiv preprint arXiv:1605.00252*, 2016.
- Har-Peled, Sariel. *Geometric approximation algorithms*, volume 173. American Mathematical Society Boston, 2011.
- Haussler, David. Decision theoretic generalizations of the pac model for neural net and other learning applications. *Information and Computation*, 100(1):78–150, 1992.
- Hoeffding, Wassily. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Koltchinskii, Vladimir. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Lecture Notes in Mathematics. Springer, 2011.
- Levrard, Clément et al. Fast rates for empirical vector quantization. *Electronic Journal of Statistics*, 7:1716–1746, 2013.
- Li, Yi, Long, Philip M, and Srinivasan, Aravind. Improved bounds on the sample complexity of learning. *Journal of Computer and System Sciences*, 62(3):516–527, 2001.
- Linder, Tamás, Lugosi, Gábor, and Zeger, Kenneth. Rates of convergence in the source coding theorem, in empirical quantizer design, and in universal lossy source coding. *IEEE Transactions on Information Theory*, 40(6):1728–1740, 1994.
- Lucic, Mario, Bachem, Olivier, and Krause, Andreas. Strong coresets for hard and soft bregman clustering with applications to exponential family mixtures. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, May 2016.
- Lucic, Mario, Faulkner, Matthew, Krause, Andreas, and Feldman, Dan. Training mixture models at scale via coresets. *To appear in Journal of Machine Learning Research (JMLR)*, 2017.
- Mendelson, Shahar. Learning without concentration for general loss functions. *arXiv preprint arXiv:1410.3192*, 2014.
- Mendelson, Shahar et al. Learning without concentration. In *International Conference on Computational Learning Theory (COLT)*, pp. 25–39, 2014.
- Moors, Johannes J A. The meaning of kurtosis: Darlington reexamined. *The American Statistician*, 40(4):283–284, 1986.
- Pollard, David. Strong consistency of k -means clustering. *The Annals of Statistics*, 9(1):135–140, 1981.
- Pollard, David. *Convergence of stochastic processes*. Springer Series in Statistics. Springer, 1984.
- Rakhlin, Alexander and Caponnetto, Andrea. Stability of k -means clustering. *Advances in Neural Information Processing Systems (NIPS)*, 19:1121, 2007.
- Ren, Yao-Feng and Liang, Han-Ying. On the best constant in Marcinkiewicz-Zygmund inequality. *Statistics & Probability Letters*, 53(3):227–233, 2001.
- Sauer, Norbert. On the density of families of sets. *Journal of Combinatorial Theory, Series A*, 13(1):145–147, 1972.
- Shamir, Ohad and Tishby, Naftali. Cluster stability for finite samples. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1297–1304, 2007.
- Shamir, Ohad and Tishby, Naftali. Model selection and stability in k -means clustering. In *International Conference on Computational Learning Theory (COLT)*, pp. 367–378. Citeseer, 2008.
- Telgarsky, Matus J and Dasgupta, Sanjoy. Moment-based uniform deviation bounds for k -means and friends. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2940–2948, 2013.
- van der Vaart, Aad W and Wellner, Jon A. *Weak convergence and empirical processes: with applications to statistics*. Springer Series in Statistics. Springer New York, 1996.
- Vapnik, Vladimir N and Chervonenkis, Alexey Ya. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.